Coordinated Science Campaign Planning for Earth Observing Missions

Robert Morris
Jennifer Dungan
Lina Khatib
Andrew Bachmann
NASA Ames Research Center

John Gasch
Paul Hempel
Jenny Williams
Terri Wood
Christopher Bruno
NASA Goddard Space Flight Center

***Abstract*— This paper addresses the problem of *coordinated Earth-science campaign planning*, the process of transforming a specification of the goals of an Earth-science campaign into a set of observations for accomplishing the campaign, utilizing diverse sensing resources from a collection of remote sensors in low Earth orbit. The paper also introduces a software architecture for a system that performs coordinated Earth science planning. The components of the architecture combine to allow for the formulation of campaign goals and plan activities, for automated or mixed-initiative (human-in-loop) plan generation and execution, and dynamic replanning. The paper also provides illustrations of the campaign planning process based on a realistic Earth science scenario requiring multiple sensing resources. This example illustrates the challenges that need to be addressed in order to generate and execute campaign plans that optimally accomplish science goals.**

I. INTRODUCTION

Science planning for satellites in low Earth orbit is currently managed independently by different mission operations centers. Coordinated science planning involving multiple sensors is done, if at all, informally among mission managers. Earth science principal investigators requiring sets of observations from different sensors have no straightforward procedure for obtaining access into mission science planning activity for the purpose of requesting time on sensors. Virtually all "coordination" of observations is accomplished on data that has already been acquired and downlinked, using graphical data archive search tools such as the EOSDIS Data Gateway (NASA), the Earth Explorer (USGS) or Space Imaging Inc.'s Carterra. Such tools provide the a single entry point into the archived data products for multiple sensors with heterogeneous capabilities.

We describe a system that would provide analogous services to Earth scientists seeking data products that have not yet been generated as the result of sensing events. The proposed system would therefore act as a portal into *science planning operations* for a set of missions. In this approach to coordinated planning, observation requests generated by an automated planner from user inputs describing campaign goals would be submitted electronically to mission operations planners, who then decide whether and how to incorporate the request into future mission schedules.

The motivation for solving coordination of data acquisition at the mission planning phase is two-fold: more effective management of sensing resources through the simultaneous deliberative planning of all the resources together; and, from the perspective of the Earth scientist, the potential for higher utility data products through the ability to more effectively control what is observed when and how.

A number of challenges, both technical and "cultural", must be addressed in developing and deploying such a system. These issues are discussed in this paper and are integrated into the design principles of the proposed system, called DESOPS (Distributed Earth Science Observation Planning System). In the next section, a realistic campaign scenario is described for purposes of motivation and illustration. There follows in section 3 a formulation of the coordinated Earth science campaign planning problem, and in section 4, a high level architectural discussion of a complete software system.

II. A CAMPAIGN SCENARIO

We use the term "*campaign*" to refer to a systematic set of activities undertaken to meet a particular science goal. Here, we present a hypothetical campaign based on a science goal to test an emissions model predicting the aerosols released by wildfires. For illustration, let us say the location of this campaign is in the southern California region, San Diego County. Data on several variables must be gathered in order to accomplish the analysis. In particular, vegetation type or biomass, atmospheric aerosol concentration and burned area are needed for the region. Fuel moisture content is a variable that also would be useful for the objectives of the science, though not a necessity.

There are several sensors that provide products at various spatial resolutions relevant to these variables. Landsat Enhanced Thematic Mapper+ (ETM+) or Thematic Mapper (TM) can be used for mapping vegetation type. Optimal timing for acquiring Landsat data for this purpose in Southern California would be June or July in the same year that the fires burned, when forested land can most easily be spectrally distinguished from grassland. For mapping aerosol concentration, images coincident to burning must be obtained. Moderate Resolution Imaging Spectroradiometer (MODIS) on the Terra and/or the Aqua satellites would provide data for this variable. MODIS data from either platform could also be used to provide coarse spatial resolution burned area after (though not too long after) the fires were out. Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) Landsat TM data would be desirable for mapping burned area with fine spatial resolution. For mapping vegetation moisture content, hyperspectral data from EO-1 Hyperion instrument are relevant. The most useful data for this purpose would need to be acquired just preceeding the fire.

This scenario, though simplified compared to what an actual scenario may involve, represents the type of integrative science currently being conducted by Earth science researchers. From this example, one can infer that inputs to a campaign planning problem consist potentially of the following characteristics:

- A set of temporal and geographic constraints on when and where images are to be taken;
- Dependencies between planned events and uncontrollable, exogenous events such as fires;
- User preferences for when an observation should be taken, or with what resource;
- A distinction between measurements that are on the critical path of analysis for meeting the science goal from others that would serve to augment the quality of the analysis, but are not strictly essential to achieving the goal.

These features combine to produce a potentially challenging problem for planning systems. In the following sections, we investigate recently developed automated planning techniques that could be applied to represent and reason about these constraints in order to generate science campaign plans.

III. Planning Problem Formulation

A coordinated Earth science campaign plan is executed by a collection of sensors. Each sensor is managed by a separate mission as part of daily mission scheduling activity [9]. Here it is assumed that missions are fundamentally "uncooperative" in the sense that each does its science planning independently of the others, with little or no direct coordination of activities. Further, individual missions are unwilling to relinquish control of the planning process for the instruments they manage; however, they are likely to accept a system that facilitates additional coordination by proposing incremental changes to their mission plans. However, decisions regarding changes to any mission schedule must be approved by the missions. This suggests a *distributed planning system* with a communication protocol whereby individual requests for observations are submitted to missions, which either accept or reject the request depending on availability of resources or other scheduling constraints.

The second fundamental feature of the planning process is the involvement of human decision-making. A *mixed-initiative* system is an intelligent system for which users input and intervention are solicited during the entire automatic reasoning process. Planning advice will take three forms [10]:

- Task advice, which allows the user to specify in detail the goals of the campaign;
- Strategic advice which recommends how the goals are to be accomplished (for example, picking a specific instrument to take an observation); and
- Evaluational advice, which specifies conditions on metrics related to the overall solution (for example, the ability to specify preferred observation windows).

Any automated solution to the campaign planning problem will need to accommodate human decision-making throughout the process.

The end-to-end planning process for Earth science campaign planning will consist of the following steps:

1) The user specifies the goals of the campaign (i.e., the set of observations and constraints involving them);
2) The system generates and displays a *flexible temporal plan* based on this input; the user adds further constraints, as desired, based on the information from the flexible plan;
3) The system enumerates and displays a subset of possible fixed plans (sequences of observation requests) that are consistent with all the constraints specified;
4) The user selects from among the list of observation requests the one(s) that are most preferred;
5) The system proceeds to execute the requested fixed plan by submitting individual requests to the relevant missions;
6) The system notifies the user of the status of the requests, which may trigger additional changes to the campaign plan.

A system for mixed-initiative plan generation and execution consists of the following core computational elements:

- An *user interface* for specifying campaign goals, which are stored in a *plan database*;
- A *planner* for generating plans based on a *constellation model* of sensors and satellite orbits; and
- A *request manager* for submitting and relaying the status of campaign requests to missions

The computational elements combine to form what will be called a *Distributed Earth Science Observation Planning System* (DESOPS), visualized in Figure 1. The remainder of this section explains how each DESOPS computational element contributes to the process of generating and executing campaign plans.

A. Specifying a campaign

A campaign request is specified as a set of observations with geographic and temporal constraints. A description of *exogenous events* that provide triggers to observation activities may also be constituent to the request. An *observation* is minimally defined in terms of the following set of attributes:

- A *type* of measurement to be taken,
- A description of a *location* on the Earth that is to be observed, and
- A *time window*, relative or absolute, within which it is to be acquired;

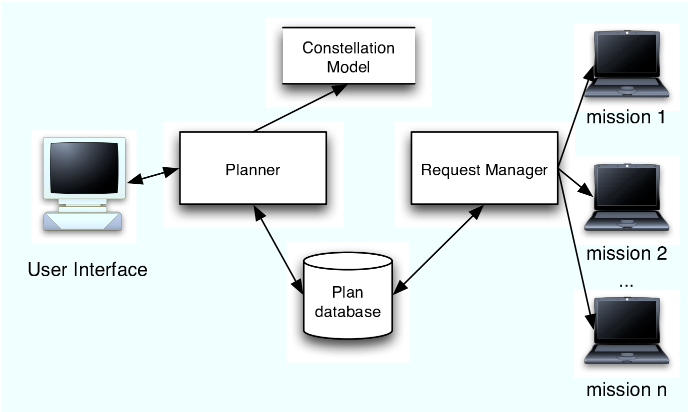

Fig. 1. Distributed Earth Science Observation Planning System

Further, a specification of the *quality* of the measurement, such as a restriction on the amount of acceptable cloud cover, may be required. Each observation attribute is associated with a domain of values, either numerical or symbolic. The elements of some of the domains can be ordered based on the specified user preferences. In particular, it is possible to impose a preference on the time of measurement, the quality of the measurement, and on other aspects of the observation such as the viewing angle for pointable instruments. For numerical domains like time, a user can apply functions that enable the expression of preferences for *minimum* or *maximum* of the values in the domain. This will allow, for example, for specifying the requirement that one measurement should be taken as soon as possible after another.

A specification of an exogenous event is required in order to formulate Earth Science campaign requirements involving Earth system occurrences such as fires, dust storms, volcanic eruptions or hurricanes. In our example, constraints arise that include observations being made a period of time before, during and after the occurrence of a large fire. An exogenous event can be specified in terms of the expected time of occurrence, or more simply as a range of times within which there is a significant probability that the event will occur. We will say that the set of observations and exogenous events together make up *activities* in a campaign.

The constraint between the onset of the dust storm and the algae bloom observation is an example of a *temporal ordering constraint*. In general, a temporal ordering constraint is a relationship between a pair of activities, where this relationship includes a time interval specifying the required gap between the activities.

Campaign data are stored in a collection of tables called the *plan database*. Each user of the system can specify one or more campaigns in a plan database. These data provide inputs into the planning process.

B. Campaign Planning

The planner transforms campaign specifications into a sequence of observation requests. There are two phases of the planning process:

- Constructing and maintaining a *flexible plan*.
- Generating sets of observation requests.

A flexible (temporal) plan is a data structure that resembles a Simple Temporal Network [3], augmented to express temporal preference information [6], as well as a means to distinguish measurement activities from exogenous events such as fires [7]. A flexible plan is so-called because it enables the representation of the permitted "slack" in scheduling times to events. This feature is useful in systems that combine planning with execution, because it allows for temporal uncertainty in the world to be explicitly represented in a plan, adding robustness during execution.

An example of an augmented flexible plan for the fire scenario is found in Figure 2. The plan is depicted as a network with a set of nodes representing the start and end points of activities. The labeled directed arcs between the nodes represent ordering constraints; for example, the arc between the node labeled "Fire end" and "Burn" expresses the constraint that the burned area observation is to be taken between 1 and 60 days after the end of the fire event, with a preference for observations taken as close to the end of the fire as possible (represented by the label *min*). There is a reference node for the beginning of the campaign plan, which is arbitrarily set to "Nov 15", the date the user initiated the campaign. There is one exogenous event, labeled "Fire", with nodes indicating its start and end. The directed arc between the specification date and the node "Fire start" is user input indicating the most likely start dates for the fire. This estimate of the start time or duration of exogenous events can be enhanced by treating the start of the fire as a random variable with an associated probability distribution. In [8], it is shown how uncertainty can be integrated into flexible plan representations. Temporal flexibility, as always, is depicted by intervals, with $[0, \infty]$ indicating that one event either happens at the same time or after another. Note that the flexible plan abstracts from considerations of which instruments are assigned to take an observation, as well as from the distinction between observations that are required to satisfy a campaign from those that are merely desirable.

The planner generates plans from a campaign specification using a *constellation model*. There are four principal components of the model: a representation of space (specifically, locations on the Earth), time, resources (specifically a collection of available sensors), and satellite orbits. Model components can be either represented as tables, or as functions or procedures that calculate values from inputs.

The minimum unit of reference for locations on the Earth consists of a single latitude/longitude coordinate pair. The simplest geometric model assumes that each lat/long specifies the center of a region of the Earth of constant proportions (e.g. the center of a region with dimensions equal to a WRS scene). More robust geometric models would contain operations for describing arbitrary regions of the Earth (for example, the ASTER scheduling system [11] contains such operations).

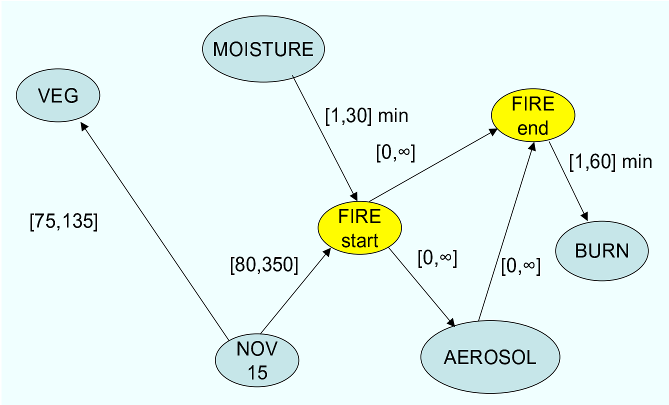

Fig. 2. A Flexible Plan for the Fire Scenario

Other models may contain non-geometric ways of specifying regions (e.g. using the names of cities).

For this model, time can be measured in discrete units of days. Temporal constraint information, as organized in a flexible plan, can be reasoned about in order to infer other constraints, or to determine whether a plan is consistent [3].

A sensor instrument has a name (e.g. ETM+), a satellite on which it resides, a type (SAR), and a specified capability, expressed in terms of the spectral, spatial or intensity parameters. In addition, it may be important to incorporate the monetary cost for acquiring an image using an instrument into the model, to enable reasoning about the relative utilities of different plans. [5]

We have viewed a campaign specification as a set of constraints on a set of observations. For each observation, there is flexibility (and, when exogenous events are involved, uncertainty) with respect to when that observation can be taken, and with what instrument. Consequently, there are potentially numerous ways of accomplishing a campaign, based on different assignments of time and instruments. Let us call each sequence of observations that accomplish a campaign a *fixed campaign*. Thus, a fixed campaign is a set of observations with specific times and instruments assigned that satisfy all the constraints in the specification. In addition to differing in time and sensor assignments, fixed campaigns will differ with respect to the degree to which optional observations are incorporated into the plan (in the fire scenario, for example, one feasible fixed campaign will contain a fuel moisture content observation, whereas another might not).

If data cost and user preferences are incorporated into the model, it is clear that there is an induced ordering of the set of fixed campaigns based on some notion of plan value or utility. The user and DESOPS planner will collectively generate the "best" plans based on this notion of utility.

C. Request Management

The *request manager* provides the interface between the planning process and the individual mission schedulers. It has two main functions: as a *plan runner* and as a mechanism for relaying the results of submitting requests from the mission to the campaign planner and user.

Inputs to the request manager will be fixed campaigns. Logically, the functionality of the request manager resembles that of a *plan runner* [7], a procedure that selects activities to execute as time passes. To be *enabled* for execution, the *active time window* of an observation must contain the current time, and any exogenous events that must precede the observation must have occurred. For example, assume that the Landsat 7 mission accepts observation requests up to 48 hours prior to scheduling a given day's observations. Then the active time window of an observation would extend from the time the observation request is generated up to 2 days before the time the observation is to be taken.

Because the DESOPS planner has limited visibility into individual mission science scheduling, there is a significant chance that observation requests might not be serviced. Consequently, it is critical to maintain a capability for *dynamic replanning* based on the results of request submissions. The trigger for rescheduling is the communication between the mission and request manager indicating the inability to service an observation request. This communication may trigger one of the following replanning activities:

- A re-submission of a request for the same measurement on the same instrument at a future time;
- An re-submission of a request for the same measurement on a different instrument at a future time;
- A "campaign abort" action;
- A revision to a campaign by adding new observations of a different type; or
- No plan revision (the campaign continues executing with no changes).

The DESOPS planner will assist the user by facilitating any of the plan revision actions initiated as a result of a mission rejecting a request.

IV. Discussion

The DESOPS system is being implemented in Java and C++, using the Automated Mission Planning and Scheduling (AMPS) system [1] as the infrastructure for building the software components and algorithms. The constraint-based approach to EOS planning and scheduling used in the design of DESOPS is based on the model formulated in [2]. The architecture presented there differs from the DESOPS architecture in adopting a centralized scheduling system for a collection of missions, rather than distributed coordinated planning. The distributed approach is preferable in not requiring significant changes to the current way of performing mission operations planning.

The DESOPS system will incorporate recent advances in automated planning. The problem of planning science observations has been addressed previously in a number of contexts. The ASPEN Planning System, developed at JPL, has been utilized for the on-board management of science activities on the Advanced Land Imager (ALI) on EO-1 [12]. Recent extensions to this work have addressed issues of coordinating science observations based on the post-processing of acquired data. Planning and scheduling of single sensors has been the subject of efforts described in [14], [13] (SPOT scheduling), [11] (ASTER scheduling), and [9] (Landsat 7 scheduling).

V. Conclusion

This paper has presented an approach to coordinated distributed planning and scheduling of Earth observations. The DESOPS concept enables mixed-initiative coordinated planning and scheduling of science observations in a distributed framework. To apply the novel approach to coordination offered here a number of technical and cultural challenges must be addressed. The core technical challenges include the following:

- **Planning with preferences and uncertainty:** Devising a effective flexible temporal planning process incorporating preferences and uncertainty;
- **Generating Optimal Plans:** Developing techniques for generating fixed plans with high expected utility incorporating user preferences and campaign costs.
- **Mixed-initiative Planner User Interface:** Developing techniques for visualizing collections of fixed plans in order to facilitate the selection of those with high utility.

References

[1] http://mopss.gsfc.nasa.gov/

[2] J. Dungan, J. Frank, A. Jonsson, R. Morris, D. Smith, Planning and Scheduling for Fleets of Earth Observing Satellites. In *Proceedings of ISAIRAS-2001*, Montreal, CA, 2001.

[3] R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks. *Artificial Intelligence*, 49, 1991:61–95, 1991.

[4] P. Pu and Denis Lalanne, Design Visual Thinking tools for Mixed Initiative Systems. *Proceedings of IUI02*, January 13-16, 2002, San Francisco, 2002.

[5] Jonsson, A. K.; Morris P. H.; Muscettola, N.; Rajan, K.; and Smith, B. 2000. Planning in Interplanetary Space: Theory and Practice. In *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling (AIPS-2000)*, 177-186.

[6] L. Khatib, P. Morris, R. Morris and F. Rossi. Temporal Reasoning about Preferences. In *Proceedings of IJCAI-01*. Morgan Kaufman, 2001.

[7] Morris, P.; Muscettola, N.; and Vidal, T. 2001. Dynamic control of plans with temporal uncertainty. In *Proceedings of the 17th International Joint Conference on A.I. (IJCAI01)*. Seattle (WA, USA): Morgan Kaufmann, San Francisco, CA.

[8] I. Tsamardinos, M. E. Pollack, and S. Ramakrishnan. Assessing the Probability of Legal Execution of Plans with Temporal Uncertainty, In *ICAPS-04 Workshop on Planning under Uncertainty and Incomplete Information*, June, 2003.

[9] W. Potter and J. Gasch. A Photo Album of Earth: Scheduling Daily Landsat 7 Activities. In *Proceedings of SpaceOps98*, Tokyo, Japan, 1998.

[10] Myers, K. L. Advisable planning systems. In Tate, A., ed. *Advanced Planning Technology*. Menlo Park, CA: AAAI Press, 1996.

[11] H. Muraoka, et. al. ASTER Observation Scheduling Algorithm. In *Proceedings of Space Ops 99*, 1998.

[12] S. Chien, et. al., Autonomous Science on the EO-1 Mission. In *Proceedings of International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, Nara, Japan, 2003.

[13] J. Agnese, N. Bataille, E. Bensana, D. Blumstein, and G. Versaillie. Exact and Approximate Methods for the Daily Management of an Earth Observation Satellite. In *Proceedings of the 5th ESA Workshop on Artificial Intelligence and Knowledge-based Systems for Space*, 1995.

[14] W. J. Wolfe, and S. E. Sorensen. Three Scheduling Algorithms Applied to the Earth Observing Systems Domain. *Management Science*, 46(1):148-168, 2000.